



IR GAME

Day After Report

Leading Through a Cyber Crisis: An Executive Workshop

Live executive tabletop exercise with an FBI cyber briefing and panel, October 13, 2026

Sponsored by: centrexIT, CCOE, Biocom, Coalition



Participants and their organizations are not named in this report. The briefing and panel speakers, as listed on Biocom's event page, are named in the Exercise Profile and described by role elsewhere, and roles are used in place of the scenario's character names. It draws on three sources: the scenario's design, points facilitators captured during the exercise, and a recorded and summarized FBI cyber briefing and panel discussion held at the same event. The company in the scenario is fictional.

Exercise Profile

Field	Detail
Exercise	Live, in-person tabletop incident response exercise for C-level and director-level leaders, followed by an FBI cyber briefing and a panel discussion
Roles in play	Executive, Technical, Medical Affairs, Legal, Finance, Business Development, Human Resources, Communications, Quality Assurance
Participant mix	Mixed leadership from a biotechnology industry audience
Briefing and panel	Josh Hohbein (centrexIT), Jess Stainer (Coalition Incident Response), Bill Reid (CCOE and National University), Vidya Murthy (MedCrypt), and Daniel Kunz (FBI Supervisory Intelligence Analyst)
Facilitation	Facilitated in person, single table

Exercise Overview

Field	Detail
Scenario Title	Leading Through a Cyber Crisis: An Executive Workshop
Date Conducted	October 13, 2026
Organization (fictional)	A clinical-stage oncology biotech of about 350 employees, with a lead candidate in trials and a partnership with a life sciences company mid-negotiation
Initial Signal	Easy to miss: an internal AI research assistant performing unusually well, a contract manufacturer fluent in internal formulation detail, and a small overnight spike in assistant traffic written off as a batch job
Attack Vector	A researcher's login captured through a personal AI tool's session tokens, then used as a valid account to query the internal AI assistant (MITRE ATT&CK T1528 and T1078)
Scoring	32 challenges across ten rounds. A perfect allocation completes 24 of 32. The best

Field	Detail
	practice path scores 510 of 702 available points (72.6%)
Total Impact	Weeks of quiet collection of formulation notes and trial design material; an anonymous tip to the partner; employees messaged directly on social media by the attacker; a \$2.5 million extortion demand against a 72-hour clock; just under \$400,000 in committed response costs; \$60 million in projected partnership value at risk; one resignation; a board asked to share the payment decision

About This Report

This report is written for everyone who took part in the event, and for anyone who wants the substance of it without having been in the room. It is deliberately not a summary of the story. The scenario was the vehicle. What follows is what the day surfaced.

Three kinds of material sit side by side. The first is the scenario's own design: the decisions it forces, the gaps it deliberately leaves in the fictional company, and the questions it puts to the room. The second is seven points the group raised during the exercise, which facilitators captured. The third is a summary of the FBI briefing and panel discussion that followed, which was recorded and condensed. Where the report describes what the fictional company did or lacked, that is the scenario's design, not a record of what any individual said. Where it describes real-world statistics and practice, it attributes them to the FBI speaker or the panel, and those figures are the speakers' own.

One framing is worth repeating before anything else. An incident response plan that has never been taken off the shelf is not a plan. It is an artifact. Everything below is aimed at the difference between the two.

The single most repeated advice of the day was to report incidents to the FBI through its Internet Crime Complaint Center (IC3) early, and if money was wired, within 24 to 48 hours.

Exercise Objectives

Derived from the scenario's learning objectives and the discussion questions the session put to participants.

- Discover whether the organization knows what its own AI tools can reach, and who would notice if that access were used by someone else.
- Find out whether the team can name who has the final word on Monday afternoon, when there is no confirmed incident, only an unexplained data pull.
- Surface whether HR, legal and security have agreed how to treat an employee whose mistake enabled an attack, before intent is known.
- Weigh containment against visibility: closing the attacker's access tells them they have been seen and ends your view of what they do next.
- Exercise communication with a partner, a board and staff while an outside party is already talking.
- Test whether the team has a view on extortion payment, sanctions checks and insurance terms before a deadline forces one.

Captured From Exercise

Seven points the group raised during the exercise. They are carried through the recommendations, actions and metrics that follow.

Have an AI policy, and train people on AI and cybersecurity. The exposure in the scenario began with a personal AI tool used for speed, not malice. A written policy says what may and may not go into outside tools, and training makes the policy something people can act on. A rule nobody has been taught is not a control.

Have tools in place to detect shadow AI. Policy alone does not tell you which teams are already using outside AI tools with company data. Detection gives you the inventory you need before an incident, rather than discovering the answer during one, as the scenario's company did on Wednesday morning.

Keep a documented incident response plan with every contact listed and defined roles, with primary and backup named. The plan in the scenario named a core cell but not who had the final word. A usable plan lists who to call, in what order, with phone numbers that work at night, and names a primary and a backup for every role so the response does not stall when someone is traveling or unreachable.

Build a relationship with the cyber insurer early. Know before an incident whether coverage is adequate, what it includes, and exactly what to do and whom to call when it is time to give notice. In the scenario a 48-hour notice clock started on Monday afternoon and nobody had settled how to meet it.

Implement phishing-resistant MFA, advanced email protection, EDR and MDR. These are the baseline controls that make a stolen or misused credential harder to use, make hostile messages harder to land, and put detection and a response team behind the endpoints. The group named them together as the minimum stack.

Follow least privilege, and separate daily-driver accounts from admin accounts. Atlas's reach was granted for convenience and never reviewed, so one researcher's login opened far more than research work needed. Least privilege shrinks what any one stolen login can reach, and separate admin accounts keep powerful credentials out of everyday browsing and email.

Agree on a communication plan beforehand covering what is told internally, to the board and externally. In the scenario every message to the partner, staff and board was written under pressure and after someone else spoke first. Deciding in advance who says what, to whom, in what order and with whose approval removes the hardest part of the week.

Captured From the FBI Briefing and Panel

The headline points from the recorded briefing and panel, in the speakers' words as summarized. The full detail is in the Practitioner Observations Catalog below.

Report to the FBI through IC3 early. The FBI investigates cyber criminals and works to impose consequences on them. It is not an incident response firm and will not tell you how to secure your network. Calling is low-key, what you share stays confidential to the extent the law allows, and the FBI does not issue press releases about victims. File even if nobody calls you back, because reports still feed its intelligence.

If money was wired, the window is 24 to 48 hours. IC3 can start a financial kill chain to try to recover wired funds. The speaker said that on average slightly over half of the money is recovered, and that after 24 to 48 hours the funds are usually gone. Speed decides the outcome.

The ransom decision belongs to your board, counsel and insurer, not the FBI. The FBI will not tell you whether to pay, but it can tell you how reliably a group has kept its promises in the past. Paying is not itself a federal crime. It becomes a serious problem if the payment goes to a sanctioned party on the Treasury Department's OFAC list. The speaker advised calling the FBI before paying, and said about 83 percent of organizations that pay are hit again by the same group.

When the target is intellectual property, nothing restores it once it leaves. Immutable backups beat encryption but not theft. State actors mostly want intellectual property, the speaker said China's

current five-year plan names biotech as a priority, and the quiet period before discovery is getting longer. Know where your crown jewels live, who can reach them and whether those systems keep logs.

Extortion has become personal. In one local case attackers sent a CEO's home address with a street-view photo and threatened violence. Some groups now call staff while posing as IT support. This is the same pattern the scenario's attacker used when messaging employees directly.

AI changes the scale of the attack. Phishing emails are now polished and fluent, voice cloning is real, and some local companies use agreed code words to confirm a caller's identity. The panel's view was that the near-term AI risk is scale, such as finding and approaching many of your employees at once, rather than autonomous hacking.

AI is exposing basics teams were told to fix 10 to 15 years ago. Know what you expose to the internet, enforce MFA and least privilege, and keep admin accounts apart from daily email accounts. AI agents receive the same access as the employee using them, so classify data, archive stale files and lock down overshared document libraries before rolling agents out.

MFA is not finished. Text-message codes are better than nothing, authenticator apps are better, and phishing-resistant MFA such as hardware keys is the goal. Session tokens can be stolen even when MFA is on, which is how the scenario's attack worked, so add managed detection and response to flag and isolate logins from unusual places or devices.

In the first hour, disconnect, do not shut down or wipe. Disconnect affected machines from the network so the evidence survives, and do not reboot or rebuild them. Backups may predate the attacker's entry and already be compromised. Bring in outside counsel, the insurer, the incident response firm and the FBI early.

Strengths Observed

What strong play looks like in this scenario, and the behaviors it is built to reward.

Escalating a signal without waiting for proof. Nothing on Monday morning is technically wrong. Strong groups still treat the manufacturer's odd fluency and the overnight traffic as worth a closer look, because by the time proof arrives the exposure has usually widened.

Naming who holds the final word before it is needed. Strong groups settle incident command early, including who decides containment, disclosure and spend, rather than letting a capable core cell drift without a decision-maker.

Treating the researcher's mistake and the main exfiltration as two problems. Fusing them too early either excuses real carelessness or accuses an employee who was also the attacker's victim. Keeping them apart until the facts connect them is the harder and better path.

Bringing in outside forensics and counsel without waiting for certainty. The ambiguity itself is the trigger. Early outside counsel also protects the investigation under privilege before sensitive findings are written down.

Protecting targeted employees within the hour. When the attacker goes around leadership and messages staff directly, the strongest response puts Human Resources and a senior leader in front of those employees immediately.

Testing the attacker's claim before reacting to it. Checking the claim of a complete development file against what the evidence shows was actually taken turns a frightening note into a decision the room can reason about.

Areas for Improvement

These describe the fictional company's position as the scenario found it. They are listed because each one has a direct counterpart in most real organizations adopting AI tools.

The AI assistant could reach far more than anyone had reviewed. Atlas's access was granted for convenience and never fully reviewed. Nobody could say what it could reach or who would notice if that scope changed.

A personal AI tool carried the credential the attacker used. Session tokens stolen from a personal AI tool let an outsider act as a real researcher. The tool was chosen because it was faster than the approved option.

The plan named a core cell but not who had the final word. Capable leads had cleared their calendars and still nobody was running the response, so containment, disclosure and spend had no owner.

Weeks of quiet collection went unseen. Detailed logs existed because extended logging was never switched off, but nothing prompted anyone to review them. An overnight spike was explained away as a scheduled job.

The partner heard about the incident from a stranger first. An anonymous contact gave the partner confidential deal details before Helix said a word, because no plan existed for what to tell a partner and when.

Insurance notice depended on details nobody had to hand. A 48-hour clock started on Monday afternoon. The broker was away and the carrier's claims line wanted a policy number and a named contact before opening a file.

Employees were not prepared to be personal targets. The attacker messaged a lab technician on a personal social media account. There was no protocol for protecting or briefing the people most likely to be approached, and one of them resigned.

Nobody had thought about extortion before the note arrived. The room had to build a position on payment, sanctions checks and insurance terms from first principles under a 72-hour clock.

Critical Decision Points

Treat an ambiguous signal as an incident, or as Monday noise. No incident has been declared. The manufacturer's fluency and the traffic spike each have innocent explanations. Someone has to decide, and decide who owns deciding.

Decide who holds the final word with no confirmed incident. The plan names a core cell but not a decision-maker. Containment, disclosure and spend all wait on the answer.

Report to the insurer inside 48 hours, or wait for a clearer picture. Reporting early risks describing something not yet understood. Waiting risks the coverage the company is most likely to need.

Close the attacker's access now, or keep watching. Closing the door stops the loss, tells the attacker they have been seen and ends the team's view of their next move. The scenario deliberately leaves this to the technical and executive players to argue out.

Approach the researcher with support or with suspicion. The same person made a real mistake and was the attacker's route in. How the first conversation happens sets the tone for the whole workforce.

Brief the partner directly, without confirming more than you know. The partner is hearing from a stranger first. A direct, accurate briefing has to be honest without outrunning the facts.

Engage an extortion specialist, or answer the demand directly. The first decision after the note arrives is not whether to pay. It is whether anyone from the company should talk to the attacker at all.

Pay the \$2.5 million demand, or refuse it inside 48 hours. The decision must account for the value of the exposed work, the partnership at risk, sanctions limits and insurance terms, and it belongs to a board that is only now being asked to share it.

Key Recommendations

Technical Controls

- Implement phishing-resistant MFA, advanced email protection, endpoint detection and response (EDR) and managed detection and response (MDR) as the baseline stack, so a misused credential or a hostile message meets controls and a response team.
- Assume session tokens can be stolen even when MFA is on. Have MDR flag logins from unusual places or devices and isolate the account, and block the Microsoft device code sign-in flow where it is not needed.
- Deploy tooling that detects shadow AI: unsanctioned AI tools, browser extensions and token grants connected to company accounts and data.
- Follow least privilege. Review and restrict what every internal AI assistant can reach to what the work actually needs, and review it on a fixed cadence.
- Before rolling out AI agents, classify data, archive stale files and lock down overshared document libraries, because an agent inherits the access of the employee using it.
- Give every administrator separate accounts for daily use and for administrative activity, with the admin account never used for email or browsing.
- Know what you expose to the internet, and patch edge devices such as VPN gateways quickly. The panel said the time from a published vulnerability to a working exploit has dropped from weeks to hours.
- Review the remote access and remote monitoring and management tools installed on your network. Ransomware actors use legitimate tools already present because they trigger fewer alerts.
- Turn on the Microsoft Purview unified audit log now, before you need it, and confirm that the systems holding your crown jewels keep logs someone reads.
- Monitor AI assistant query patterns for unusual volume, timing and data categories, and alert on a valid account behaving unlike its owner.
- Keep immutable backups, and treat them as protection against encryption only. They do not help once research data has been stolen.

Policy and Procedures

- Publish an AI policy that says what may and may not be entered into outside tools, names an approved alternative that is as fast as what people were using, and says who to ask.
- Keep a documented incident response plan that lists every contact to call, with working numbers, and defines each role with a named primary and backup. Keep a laminated who-to-call list covering the insurer, outside counsel, the incident response firm, key vendors and the FBI.
- Name one incident commander with backups, and write down who has the final word on containment, disclosure and spend. The lead role may shift from recovery to forensics as the response moves along.
- Agree on a communication plan in advance for what is told internally, to the board and externally, with one approval path.

- Write a business continuity plan that says which systems must stay up and how you would operate manually, with leadership, including the C-suite, helping to build the playbooks and naming a delegate if the decision-maker cannot be reached.
- Set a board notification threshold tied to the materiality of an exposure, not to the next scheduled meeting.
- Take a written position on extortion before a demand arrives, including who decides, what proof is required, an OFAC sanctions check, and a call to the FBI before any payment.
- Build an employee protection protocol for staff approached personally by an attacker, including a code word or callback routine for urgent requests that claim to come from leadership or IT support.
- Include cybersecurity terms in vendor and device contracts, using ready-made model contract language where it exists.

Training and Awareness

- Train the whole workforce on AI and cybersecurity together, including why personal AI tools are a risk and how to ask for an approved one.
- Teach staff how to report AI behavior that seems unusually good as well as behavior that is obviously broken.
- Brief leadership on the extortion decision before they face it. The call will be theirs and few have rehearsed it.
- Prepare staff for fluent AI-written phishing, cloned voices and callers posing as IT support, and tell them how to verify and whom to tell.
- Prepare staff to be approached on personal accounts, and tell them what to do and whom to tell.
- Run tabletop exercises at least twice a year with legal, communications, human resources and executive leaders. The panel noted that holidays, long weekends and Friday afternoons are when incidents tend to hit.

Business Process Improvements

- Build a relationship with the cyber insurer early. Review the policy with a knowledgeable third party, learn what it covers and excludes, and record exactly whom to call and what the carrier needs in order to open a file.
- Check that the policy covers business interruption, including an outage at a key vendor.
- Confirm that a policy number, broker contact and named company contact are always to hand, including when the broker is away.
- Pre-select outside counsel, a forensic firm and an extortion response specialist, and confirm current rates rather than assuming the old retainer holds. Counsel can often file the IC3 report for you.
- Establish a working relationship with the local FBI field office, and have a staff member join InfraGard to receive indicators of compromise and threat reports.
- Establish emergency spending authority so leaders can commit outside help without waiting for a budget cycle.
- Know your crown jewels, put a value on them, and protect them first. Aim for higher framework tiers such as CIS Controls Implementation Group 3 and budget for them.
- Buy software only from vetted, reputable vendors and make someone responsible for confirming known vulnerabilities are patched.

If You're Hit: Who to Call

From the briefing and panel. Put this in your incident response plan now, with a named owner for each call.

- 1. Lock down the money.** If funds were wired, report to IC3 and your bank right away. The recovery window is about 24 to 48 hours.
- 2. Bring in your partners.** Call outside counsel, your cyber insurer and your incident response firm. Counsel can often file the IC3 report for you.
- 3. Call the FBI.** Use your local field office or IC3. Call before deciding on any ransom payment, and sooner still if you suspect a nation-state actor, since some intelligence can be shared with victims once the FBI knows who you are.
- 4. Preserve evidence.** Disconnect affected systems from the network. Do not shut them down, reboot them or wipe them.

Cleaning Up a Compromised Microsoft 365 Account

A checklist from the panel. Business email compromise was the attack the panel said it sees most.

- ☐ Reset the password and revoke all active sessions.
- ☐ Remove all MFA methods, wait about 30 minutes, then re-register them, since attackers often add their own.
- ☐ Check for suspicious inbox rules and new enterprise applications. Attackers use legitimate apps to copy whole mailboxes.
- ☐ Review SharePoint and OneDrive activity.
- ☐ Turn on the Microsoft Purview unified audit log before you need it. The panel said it is available at any license level. Per Microsoft, it can take several hours to start collecting and keeps 180 days of data by default.

Follow-up Actions

Immediate, 0 to 30 days

- Find out what your internal AI tools can reach today and who would know if that changed.
- Open a no-blame window for teams to disclose personal AI tool use with company data.
- Check that your incident response plan lists every contact with a working number, and that each role has a named primary and backup. Print the who-to-call list.
- Find your cyber insurance policy, confirm the notice window, the business interruption terms and the exact steps and contacts for giving notice, and book time with your broker or insurer.
- Turn on the Microsoft Purview unified audit log if it is not already on.
- Review what remote access and remote monitoring tools are installed on your network.
- Write down who has the final word during an incident, and circulate it.

Short-term, 30 to 90 days

- Publish an AI policy and start AI and cybersecurity training for all staff, naming an approved tool.
- Deploy or verify shadow AI detection and decide who reviews what it finds.
- Confirm phishing-resistant MFA, advanced email protection, EDR and MDR coverage, and close the gaps. Block the device code sign-in flow where it is not needed.
- Separate daily-driver and administrator accounts, and remove standing admin rights where they are not needed.
- Classify data, archive stale files and lock down overshared libraries before any AI agent rollout.

- Draft the communication plan for internal, board and external audiences, with holding statements and an approval path.
- Write the extortion position, including the OFAC check, insurance terms and a call to the FBI before payment.
- Have a staff member join InfraGard and make contact with your local FBI field office.

Long-term, 90 days and beyond

- Write the business continuity plan: which systems must stay up, how you operate manually, and who holds authority with named delegates.
- Review every AI tool and integration against least privilege on a recurring schedule.
- Identify your crown jewels, confirm the systems holding them keep logs, and work toward a higher framework tier.
- Exercise the incident response plan and the communication plan together at least twice a year, with a backup person standing in for each primary role at least once.
- Fold what this scenario showed about AI tools, integrations and trusted credentials into the incident response plan.

Metrics and Success Criteria

Metric	Target
AI policy published, with AI and cybersecurity training completed by staff	Policy in place; 90 percent of staff trained within 12 months, refreshed annually
Shadow AI detection coverage	Detection in place across managed devices and accounts, reviewed monthly
Incident response plan contact list verified	Every contact tested at least twice a year
Incident response roles with named primary and backup	100 percent of roles
Cyber insurance policy reviewed, with notice steps and business interruption terms documented	Reviewed annually with a third party; notice steps rehearsed
Phishing-resistant MFA coverage	100 percent of privileged and remote access accounts
EDR and MDR coverage	100 percent of managed endpoints, with response tested
Admin activity on separate accounts	100 percent of administrators, verified quarterly
Unified audit logging enabled, and logging confirmed on crown jewel systems	100 percent of in-scope systems, verified quarterly
Review of installed remote access and monitoring tools	At least quarterly, with unapproved tools removed
AI tool access-scope review	At least quarterly, documented
Time to report a suspected wire fraud to the bank and IC3	Within 24 hours, since recovery odds fall sharply after 24 to 48 hours
Time to disconnect affected systems rather than shut them down	Within 1 hour of confirmation, with evidence preserved
Communication plan covering internal, board and external audiences	Written, approved and exercised at least annually
Time from anomaly to a named incident decision-maker	Under 4 hours
Incident response exercises conducted per year	At least two, with executive leadership and legal

Regulatory and Framework Mapping

Regulatory obligations vary by jurisdiction and sector. Confirm the specific requirements that apply to your organization with counsel rather than relying on this table.

Incident Aspect	Framework or Regulation	Gap Identified	Recommended Action
Unreviewed AI assistant access	CIS Controls 3.3 and 6.8	Access granted for convenience and never reviewed	Restrict to role-based access and review on a fixed cadence
Personal AI tools with proprietary data	CIS Controls 2.3 and 14.4	No approved, equally fast alternative; no training on data handling	Publish an AI policy, name an approved tool and train staff
Stolen tokens acting as a valid account	CIS Controls 6.2 and 5.3; MITRE ATT&CK T1528, T1078	Tokens outside company control could reach company systems	Phishing-resistant MFA, token revocation, MDR, shadow AI detection
No named decision-maker	CIS Controls 17.1 and 17.5	Plan names a cell but not who has the final word	Name primary and backup for every role and one incident commander
Incident messaging under pressure	CIS Control 17.6	No single approval path for partner, staff, board and media messages	Agree on the communication plan beforehand
Cyber insurance notice and coverage	Policy terms (contractual)	48 hours from confirmation; broker away; details needed to open a file; business interruption terms unchecked	Relationship with the insurer early; documented notice steps; confirm business interruption cover
Ransom payment	OFAC sanctions guidance	No sanctions review before a payment decision	Complete a sanctions and legal review, and call the FBI before any payment
Cybercrime reporting	FBI and IC3 (voluntary)	No agreed trigger to report; wire recovery window is 24 to 48 hours	Decide in advance who reports; counsel can often file; report as early as practical
Trade secret theft	Defend Trade Secrets Act	Evidence and privilege at risk during containment	Early outside counsel; preserve logs and disconnect rather than wipe
Evidence and audit logging	CIS Controls 8.2, 8.10 and 8.11	Logging not enabled, or not reviewed, on systems holding crown jewels	Enable unified audit logging and confirm crown jewel systems keep reviewed logs
Regulator information sharing	FDA and SEC information-sharing agreement, August 2026	Scrutiny of public biotech and device companies is rising	Have counsel review how an incident and its disclosure would read to both agencies
Securities disclosure	SEC Form 8-K Item 1.05, if publicly traded	Four business days after a materiality determination	Have counsel set the materiality process in advance
Breach notification laws (personal data)	State and sector breach laws	Not triggered as scripted; exposed data is proprietary, not personal	Confirm with counsel which obligations apply to non-personal data

Lessons Learned

What Worked Well

Some preparation paid off before it was needed. Extended logging that was never switched off, access lists already on file, an outside counsel hotline and a standing spending authority each gave the response a head start. None were built for this incident.

Testing the claim before reacting to it. Comparing the attacker's claim against reconstructed data loss keeps the room reasoning rather than reacting.

Treating the payment question as a business decision with legal limits. Sanctions, insurance terms and the cost of the partnership give the decision something firmer than instinct, and the FBI speaker confirmed it belongs to the board, counsel and insurer.

What Did Not Work

A plan existed without a decision-maker. Naming a core cell was not enough. Without a named final word, capable people waited on each other.

Everything outward-facing was reactive. The partner, staff, forums and a reporter all moved before the company spoke. The cause is that no one had decided in advance what to say, to whom and when.

AI tools were governed by convenience. The assistant's reach and the personal tools' reach were both chosen for speed, so a single login opened far more than the work required.

Detection depended on luck. The logs were there. Weeks passed because nothing made anyone read them.

The FBI was not in the plan until Thursday. The scenario opens law enforcement reporting only on Thursday. The speaker's advice was to report early, and sooner still when a nation-state actor is suspected.

Unexpected Findings

Closing the door has a cost. Shutting the attacker out ends the team's view of what they do next. Many participants had never weighed containment against visibility.

The suspected partner was innocent. The manufacturer's knowledge had an ordinary explanation: a routine technical transfer package that was supposed to go. Time spent suspecting them was time lost.

The FBI will not tell you whether to pay. It can describe how reliably a group has kept its promises, but the decision belongs to the board, counsel and insurer. Several participants had expected a recommendation.

Paying is legal until it is not. The speaker said paying a ransom is not a federal crime. The serious problem is payment to a sanctioned party, which is why the OFAC check comes first.

Payment often brings a second demand. Per the speaker, about 83 percent of organizations that pay are hit again by the same group, which matches the scenario's warning that a single payment is not a resolution.

MFA does not stop a stolen session. The panel confirmed that session tokens can be stolen even with MFA on, which is exactly how the scenario's attacker acted as a valid employee.

Personal social media became the attack surface. The pressure reached leadership through a lab technician's personal account, a place no company control reaches. The speaker described similar escalation against a real CEO's home.

Proprietary data is outside the usual breach playbook. The exposed material was not personal data, so the notification assumptions most teams carry did not apply, and the obligations sat elsewhere.

ACTIONS

Category	Action
Review	Review your cyber insurance policy and your relationship with the carrier: what is covered, whether business interruption includes a vendor outage, who to call and what they need to open a file.
Investigate	Investigate whether shadow AI use exists today, and which remote access tools are installed on your network.
Coordinate	Coordinate your communication plan for internal, board and external audiences, and make contact with your local FBI field office and InfraGard before you need them.
Create	Create an AI policy and a documented incident response plan with every contact listed, a laminated who-to-call list, and a primary and backup for each role.
Train	Train staff on AI and cybersecurity together, including AI phishing and voice cloning, and brief leaders on the extortion decision.
Implement	Implement phishing-resistant MFA, advanced email protection, EDR and MDR, least privilege, separate admin accounts, and unified audit logging.
Check	Check what your internal AI tools can reach, where your crown jewels live, and whether those systems keep logs.
Schedule	Schedule the next exercise, at least twice a year, with legal, communications, HR and executives at the table.

Practitioner Observations Catalog

Every useful observation, technique and piece of information from the FBI briefing and panel, given a title and a short description. It is drawn from a recorded and summarized account of the session rather than a verbatim transcript. Statistics are the speakers' own and are attributed to them, and some may be rounded or summarized. Most of what follows came from practitioners describing what they have actually seen.

What the FBI does, and what it does not

- 1. The FBI investigates cyber criminals; it is not an incident response firm.** It works to impose consequences on criminals and will not tell you how to secure your network. You still need your own response firm.
- 2. Almost every current case has international ties.** The FBI works with partner agencies and with legal attachés in U.S. embassies abroad.
- 3. Calling the FBI is low-key.** Agents arrive without fanfare, what you share stays confidential to the extent the law allows, and there are no press releases about victims.
- 4. Report through IC3, and report even if nobody calls back.** Analysts route each complaint to the right field office, and reports still feed the FBI's intelligence.
- 5. Counsel can often file the IC3 report for you.** The panel suggested bringing in outside counsel, the insurer, the response firm and the FBI early, with counsel filing where possible.
- 6. The FBI sometimes holds decryption keys.** When it had infiltrated the Hive ransomware group, the speaker said it gave victims the key at no cost. Asking the FBI early can matter.

Reporting and recovering wired money

7. IC3 can start a financial kill chain. If funds were wired, IC3 can try to recover them. The speaker said slightly over half of the money is recovered on average.

8. Speed decides recovery. After 24 to 48 hours the funds are usually gone, so the call to the bank and IC3 cannot wait for the facts to settle.

The ransom decision

9. The FBI will not tell you whether to pay. That decision belongs to your board, counsel and insurer.

10. The FBI can describe a group's track record. It can tell you how reliably a group has kept its promises in the past, which is useful input to the decision.

11. Paying is not a federal crime; paying a sanctioned party is a serious problem. State-linked groups in North Korea or Iran on the OFAC list are the example given. Check before you pay.

12. Call the FBI before you pay. The speaker was explicit that the call comes before payment, not after.

13. Expect to be hit again. The speaker said about 83 percent of organizations that pay are hit again by the same group.

Trends to watch

14. Complaint volume is enormous and rising. IC3 received more than 1 million complaints in 2025, with about \$20.9 billion in reported U.S. losses, and California ranked first in both complaints and losses.

15. Business email compromise is expensive and underreported. California losses topped \$430 million in 2025, with about \$160 million tied to AI by the speaker. Private research suggests actual losses are about 20 times what is reported.

16. Attackers live off the land. Ransomware actors use legitimate tools already on your network, including remote access and remote monitoring and management tools such as ScreenConnect, TeamViewer and Action1, which are less likely to trigger alerts. Review what software is installed.

17. Extortion is escalating to the personal. In one local case attackers sent a CEO's home address with a street-view photo and threatened violence.

18. Some groups call staff posing as IT support. The FBI has warned about this, mainly against law firms. Staff need a way to verify a caller.

19. Social engineering is still the top way in. AI makes phishing emails polished and fluent, removing the clumsy wording people used to rely on.

20. Voice cloning is real. Attackers impersonate a CEO on the phone to push urgent wire transfers. Audio-only fakes are harder to spot than video, and some local companies now use agreed code words.

21. The near-term AI risk is scale. The speaker's example was a model finding and approaching many of your employees at once, rather than an autonomous attacker.

22. Carelessness is an insider threat too. A contractor who never patches a known vulnerability counts. Buy software only from vetted, reputable vendors.

Nation-state actors and intellectual property

23. State actors mostly want intellectual property. The speaker said China's current five-year plan names biotech as a priority.

24. They stay hidden, and the hidden period is lengthening. Nation-state actors stay quiet to protect their zero-day exploits, so discovery tends to come late.

25. Know your crown jewels and protect them first. Know where they live, who can access them and whether those systems keep logs. Many teams find out mid-incident that a costly platform has no logging.

26. Iran is a growing threat. This includes state hackers and independent sympathizers. The speaker cited the recent Stryker compromise and attacks on U.S. water facilities.

27. If you suspect a nation-state actor, call the FBI even sooner. Some of its intelligence can be shared with victims, but only once it knows who you are.

28. Join InfraGard to receive indicators of compromise. It is an FBI-led public and private partnership, and members get threat reports and quarterly meetings.

29. Immutable backups beat encryption, not theft. Once research data leaves, you cannot restore it. This is why the scenario's loss is permanent whatever happens to the ransom.

Basics, AI and agents

30. AI is exposing basics teams were told to fix 10 to 15 years ago. The panel's central point was that the gaps are old ones, now exploited faster.

31. Start by knowing what you expose to the internet. Then enforce MFA and least privilege, and keep admin accounts separate from daily email accounts.

32. Write the policies. They are dull, but the panel said they decide how well you respond.

33. AI agents get the same access as the employee using them. Classify your data, archive stale files and lock down overshared SharePoint before you roll agents out.

34. Take new risks on purpose. Connected medical devices and fast AI adoption are growing the attack surface faster than regulation is keeping up, so accept those risks deliberately.

35. Edge devices are the target, and the clock has collapsed. Attackers are going after devices such as VPN gateways. With AI, the time from a published vulnerability to a working exploit has dropped from weeks to hours.

MFA and identity

36. Move up the MFA ladder. Text-message codes are better than nothing, authenticator apps such as Microsoft Authenticator or Duo are better, and phishing-resistant MFA such as hardware security keys is the goal.

37. Session tokens can be stolen even when MFA is on. Business email compromise was the attack the panel sees most. Add managed detection and response so unusual logins are flagged and the account isolated.

38. Watch for device code phishing. Attackers abuse Microsoft's real sign-in code flow, the same enter-this-code step used to log a TV into a streaming app, to take over a session. You never see a fake page, so block the flow where you can.

The first hour of a suspected breach

39. Name roles with a primary and a backup for each. Keep a laminated who-to-call list covering your insurer, outside counsel, response firm, key vendors and the FBI.

40. Practice with tabletop exercises. The panel noted that holidays, long weekends and Friday afternoons are when incidents tend to hit.

41. Bring in counsel, the insurer, the response firm and the FBI early. Do not wait until you understand the incident.

42. Disconnect; do not shut down or reboot. Disconnecting affected machines from the network preserves the evidence.

43. Do not rush to wipe and restore. You may destroy evidence, and your backups may predate the attacker's entry and already be compromised.

44. Get to know your cyber insurance contact before you need them. The relationship matters most when the notice clock is running.

Cleaning up a compromised email account

45. Reset the password and revoke all sessions. A password reset alone leaves stolen sessions alive.

46. Remove all MFA methods, wait about 30 minutes, then re-register. Attackers often add their own authentication methods.

47. Check inbox rules and new enterprise applications. Attackers use legitimate apps to copy whole mailboxes.

48. Review SharePoint and OneDrive activity. The compromise rarely stops at the mailbox.

49. Turn on the unified audit log now. The panel said it is available at any license level. Per Microsoft, it can take several hours to start collecting and keeps 180 days of data by default.

Keeping research and patient services running

50. Decide ahead of time which systems must stay up. A business continuity plan should also say how you would operate manually.

51. Leadership should help build the playbooks. The C-suite should agree who has authority and name a delegate in case that person cannot be reached.

52. Check that insurance covers business interruption. That includes outages at a key vendor.

53. Name one incident commander, with backups. The lead role may shift from recovery to forensics as the response moves along.

54. Aim for higher framework tiers and budget for them. The panel pointed to CIS Controls Implementation Group 3 for high-risk organizations.

55. Regulatory scrutiny is rising. The panel said an August 2026 agreement lets the FDA share nonpublic information about public device and biotech companies with the SEC.

Your Action Commitments

These are the three questions the exercise closes on. They are left open here deliberately. A report that answers them for you is a document; one you answer yourself is a commitment.

What would an incident like this cost your organization?

What are two or three actions you are taking back to your team?

What would you do differently in your own organization?

Conclusion

The scenario ends without resolution, and that is deliberate. By Friday afternoon the demand is outstanding with less than 48 hours left, a valued employee has resigned, and a board is being asked to share a decision nobody in the room has made before. Real incidents rarely offer a clean finish, and an exercise that provides one teaches the wrong lesson.

The failures at the center of the story were not mainly technical. A stolen login was the delivery mechanism. The rest was an assistant with unreviewed reach, a personal shortcut nobody could see, a plan that named a team but not a decision-maker, an insurer nobody had spoken to before the clock started, messages written only after someone else spoke first, and a payment decision that arrived with no position behind it. The FBI briefing and panel said the same thing from the real world: AI is exposing basics teams were told to fix a decade ago, session tokens defeat MFA that is not phishing-resistant, and the first hour is decided by whether you have a plan with named roles, a who-to-call list and the discipline to disconnect rather than wipe.

The group's own conclusions point the same way: an AI policy and training, tools to detect shadow AI, a documented plan with contacts and backups, an early relationship with the insurer, phishing-resistant MFA with email protection, EDR and MDR, least privilege with separate admin accounts, and a communication plan agreed beforehand. Almost all of it is addressable before the next incident.

If this day produces one change, make it the smallest one with the largest reach: write down who has the final word and who calls the FBI, name a backup for each, and tell people. Most of the rest becomes easier once someone is permitted to say that the thing has started.

Resources and References

The resources below were mentioned by the FBI speaker or the panel, as listed in the event recap. A short list of additional references for the AI themes follows, and those were not checked from this session.

Mentioned at the briefing and panel

IC3, Internet Crime Complaint Center

<https://www.ic3.gov>

File a cybercrime report with the FBI. This is also how a wire-fraud recovery, the financial kill chain, gets started. Mentioned by the FBI speaker and the panel.

FBI San Diego Field Office

<https://www.fbi.gov/sandiego>

Local FBI contact for San Diego and Imperial counties, reachable around the clock at (858) 320-1800 or at <https://tips.fbi.gov>. Mentioned by the FBI speaker. Use your own local field office if you are elsewhere.

2025 IC3 Internet Crime Report

https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf

Source for the briefing's trend data: 1,008,597 complaints and about \$20.9 billion in reported U.S. losses in 2025. Mentioned by the FBI speaker.

InfraGard

<https://infragard.fbi.gov>

FBI-affiliated public and private partnership. Members get threat reports, indicators of compromise and quarterly meetings. Mentioned by the FBI speaker and the panel.

InfraGard chapters

<https://www.infragardnational.org/about-us/infragard-chapters/>

Find your local chapter. The San Diego chapter covers San Diego and Imperial counties.

OFAC Sanctions List Search

<https://sanctionssearch.ofac.treas.gov/>

Treasury Department tool for checking whether a person or group is sanctioned. Relevant to any ransom payment decision. Mentioned by the FBI speaker.

DOJ: Hive ransomware disruption

<https://www.justice.gov/archives/opa/pr/us-department-justice-disrupts-hive-ransomware-variant>

The case the FBI speaker described: the FBI handed out more than 1,300 decryption keys and averted about \$130 million in ransom payments.

FBI alert coverage: fake IT-support calls (Silent Ransom Group)

<https://www.bleepingcomputer.com/news/security/fbi-warns-of-silent-ransom-group-in-person-data-theft-attacks/>

The FBI warning about attackers posing as IT support, mainly against law firms. Mentioned by the FBI speaker.

GovRAMP (formerly StateRAMP)

<https://govramp.org>

Security authorization for cloud vendors selling to state and local government. Mentioned by the FBI speaker.

San Diego Cyber Clinic (CCOE)

<https://sdccoe.org/cyber-clinic/>

Free cybersecurity help for small businesses and under-resourced organizations, from policy writing to vulnerability assessments, delivered by students from National University, CSU San Marcos and SDSU. Mentioned by the panel.

IRGame

<https://www.irgame.ai/>

Online incident response tabletop game. Free public sessions run the last Friday of every month. Mentioned by the panel.

CISA: Implementing Phishing-Resistant MFA

<https://www.cisa.gov/sites/default/files/publications/fact-sheet-implementing-phishing-resistant-mfa-508c.pdf>

Why SMS and app codes can be phished, and how to roll out stronger MFA in phases. Mentioned by the panel.

Microsoft: Block device code flow

<https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-authentication-flows>

Microsoft's guidance on blocking the sign-in flow abused in device code phishing. Mentioned by the panel.

Microsoft Purview: Turn on audit logging

<https://learn.microsoft.com/en-us/purview/audit-log-enable-disable>

How to enable the Microsoft 365 unified audit log. Microsoft says records are kept for 180 days by default. Mentioned by the panel.

NIST Cybersecurity Framework 2.0

<https://www.nist.gov/cyberframework>

A starting framework for building a security program. Mentioned by the panel.

CIS Critical Security Controls, Implementation Groups

<https://www.cisecurity.org/controls/cis-controls-implementation-groups>

Prioritized controls in tiers, from IG1, basic hygiene, to IG3, high-risk organizations. Mentioned by the panel.

CISA Incident and Vulnerability Response Playbooks

<https://www.cisa.gov/resources-tools/resources/federal-government-cybersecurity-incident-and-vulnerability-response-playbooks>

Step-by-step response checklists that any organization can adapt. Mentioned by the panel.

HSCC Model Contract Language for Medtech Cybersecurity v2

<https://healthsectorcouncil.org/wp-content/uploads/2025/11/MC2v2.pdf>

Ready-made contract language for cybersecurity terms between device makers and the organizations that buy and use them. Mentioned by the panel.

HSCC publications

<https://healthsectorcouncil.org/hsc-cc-publications/>

Further publications from the Health Sector Coordinating Council, listed alongside the model contract language.

SEC and FDA information-sharing agreement (summary)

<https://corpgov.law.harvard.edu/2026/09/22/fda-and-sec-open-a-new-information-sharing-channel-implications-for-public-life-sciences-companies/>

Summary of the August 31, 2026 agreement letting the FDA share nonpublic information about regulated companies with the SEC. Mentioned by the panel.

Additional references for the AI themes

Not from the session.

NIST AI Risk Management Framework

<https://www.nist.gov/itl/ai-risk-management-framework>

A foundation for governing AI systems, relevant to scoping and reviewing what an internal assistant can reach.

CISA and NSA, Deploying AI Systems Securely

<https://www.cisa.gov/resources-tools/resources/deploying-ai-systems-securely>

Joint guidance on securing AI deployments, covering the access and integration gaps in this scenario.

OWASP Top 10 for Large Language Model Applications

<https://owasp.org/www-project-top-10-for-large-language-model-applications/>

Risks specific to LLM-connected tools, including excessive agency and insecure integrations.

Where to Go From Here

The exercise was built to surface gaps, not to grade anyone. If the day raised questions about your own plan, your AI tools or your insurance terms, the useful next step is a short conversation while the details are fresh.

Book a 30-minute cybersecurity assessment. A 30-minute session with the centrexIT security team that walks through your incident response plan, your who-to-call list and your baseline controls (MFA, email protection, EDR and MDR) against what the day surfaced. You leave with a short list of what to fix first.

Schedule your 30-minute cybersecurity assessment

Run this exercise with your own leadership team. The tabletop you played can be run privately for your own executives, using your organization's structure, your real roles and your own plan. It is the most direct way to find out who holds the final word before you need to know.

Questions about anything in this report? Reply to the email this report came with, or contact Josh Hohbein, who facilitated the exercise, at jhohbein@centrexit.com.

About centrexIT. centrexIT is a San Diego-based managed IT services and cybersecurity company, operating since 2002, with a team of about 60 people. We help organizations adopt AI safely because we already understand their people, systems, data, security posture and operational risk. centrexIT is a Biocom Endorsed Partner. People-First. AI-Amplified. Cyber IT.